

Frugal Mechanism Design via Spectral Techniques

Ning Chen*

Edith Elkind*

Nick Gravin^{*†}

Fedor Petrov[†]

Abstract

We study the design of truthful mechanisms for set systems, i.e., scenarios where a customer needs to hire a team of agents to perform a complex task. In this setting, frugality [2] provides a measure to evaluate the “cost of truthfulness”, that is, the overpayment of a truthful mechanism relative to the “fair” payment.

We propose a uniform scheme for designing frugal truthful mechanisms for general set systems. Our scheme is based on scaling the agents’ bids using the eigenvector of a matrix that encodes the interdependencies between the agents. We demonstrate that the r -out-of- k -system mechanism and the $\sqrt{\cdot}$ -mechanism for buying a path in a graph [21] can be viewed as instantiations of our scheme. We then apply our scheme to two other classes of set systems, namely, vertex cover systems [6, 10] and k -path systems, in which a customer needs to purchase k edge-disjoint source-sink paths. For both settings, we bound the frugality of our mechanism in terms of the largest eigenvalue of the respective interdependency matrix.

We show that our mechanism is optimal for a large subclass of vertex cover systems satisfying a simple local sparsity condition. For k -path systems, while our mechanism is within a factor of $k + 1$ from optimal, we show that it is, in fact, *optimal*, when one uses a modified definition of frugality proposed in [10]. Our lower bound argument combines spectral techniques and Young’s inequality, and is applicable to all set systems. As both r -out-of- k systems and single path systems can be viewed as special cases of k -path systems, our result improves the lower bounds of [21] and answers several open questions proposed in [21].

*Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore. Email: ningc@ntu.edu.sg, eelkind@ntu.edu.sg, ngravin@pmail.ntu.edu.sg.

[†]St.Petersburg Department of Steklov Mathematical Institute RAS, Russia.

1 Introduction

Consider a scenario where a customer wishes to purchase the rights to have data routed on his behalf from a source s to a destination t in a network where each edge is owned by a selfishly motivated agent. Each agent incurs a privately known cost if the data is routed through his edge, and wants to be compensated for this cost, and, if possible, make a profit. The customer needs to decide which edges to buy, and wants to minimize his total expense.

This problem is a special case of the *hiring-a-team* problem [27, 21, 20, 5, 10]: Given a set of agents \mathcal{E} , a customer wishes to hire a team of agents capable of performing a certain complex task on his behalf. A subset $S \subseteq \mathcal{E}$ is said to be *feasible* if the agents in S can jointly perform the complex task. This scenario can be described by a *set system* $(\mathcal{E}, \mathcal{F})$, where \mathcal{E} is the set of agents and \mathcal{F} is the collection of feasible sets. Each agent $e \in \mathcal{E}$ can perform a simple task at a privately known cost $c(e)$. In such environments, a natural way to make the hiring decisions is by means of *mechanisms* — Each agent e submits a *bid* $b(e)$, i.e., the payment that he wants to receive, and based on these bids the customer selects a feasible set $S \in \mathcal{F}$ (the set of *winners*), and determines the payment to each agent in S .

A desirable property of mechanisms is that of *truthfulness*: It should be in the best interest of every agent e to bid his true cost, i.e. to set $b(e) = c(e)$ no matter what bids other agents submit; that is, truth-telling should be a dominant strategy for every agent. Truthfulness is a strong and very appealing concept: it obviates the need for agents to perform complex strategic computations, even if they do not know the costs and strategies of others. This property is especially important in the Internet and electronic commerce settings, as most protocols are executed instantly.

One of the most celebrated truthful designs is the VCG mechanism [28, 7, 18], where the feasible set with the smallest total bid wins, and the payment to each agent e in the winning set is his threshold bid, i.e., the highest value that e could have bid to still be part of a winning set. While VCG mechanism is truthful, on the negative side, it can make the customer pay far more than the true cost of the winning set, or even the cheapest alternative, as illustrated by the following example: There are two parallel paths P_1 and P_2 from s to t , P_1 has one edge with cost 1 and P_2 has n edges with cost 0 each. VCG selects P_2 as the winning path and pays 1 to every edge in P_2 . Hence, the total payment of VCG is n , the number of edges in P_2 , which is far more than the total cost of both P_1 and P_2 .

The VCG overpayment property illustrated above is clearly undesirable from the customer's perspective, and thus motivates the search for truthful mechanisms that are *frugal*, i.e., select a feasible set and induce truthful cost revelation without resulting in high overpayment. However, formalizing the notion of frugality is a challenging problem, as it is not immediately clear what the payment of a mechanism should be compared to. A natural candidate for this benchmark is the total cost of the closest competitor, i.e., the cost of the cheapest feasible set among those that are disjoint from the winning set. This definition coincides with the second highest bid in single-item auctions and has been used in, e.g., [1, 2, 27, 11]. However, as observed by Karlin, Kempe and Tamir [21], such feasible set may not exist at all, even in monopoly-free set systems (i.e., set systems where no agent appears in all feasible sets). To deal with this problem, [21] proposed an alternative benchmark, which is bounded for any monopoly-free set system and is closely related to the buyer-optimal Nash equilibrium of first-price auctions (see Definition 2.2). Nash equilibrium corresponds to a stable outcome of the bargaining process, and therefore provides a natural lower bound on the total payment of any dominant strategy mechanism. Throughout the paper, we use the benchmark of [21], as well as its somewhat more relaxed variant suggested in [10] to study frugality of truthful mechanisms.

1.1 Our Results

Uniform Frugal Truthful Mechanisms. We propose a uniform scheme, which we call PRUNING-LIFTING MECHANISM, to design frugal truthful mechanisms for set systems. At a high-level view, this mechanism consists of two key steps: pruning and lifting.

- **Pruning.** In a general set system, the relationships among the agents can be arbitrarily complicated. Thus, in the pruning step, we remove agents from the system so as to expose the structure of the competition. Intuitively, the goal is to keep only the agents who are going to play a role in determining the bids in Nash equilibrium; this enables us to compare the payoffs of our mechanism to the total equilibrium payment. Since we decide which agents to prune based on their bids, we have to make our choices carefully so as to preserve truthfulness.
- **Lifting.** The goal of the lifting process is to “lift” the bid of each remaining agent so as to take into account the size of each feasible set. For this purpose, we use a graph-theoretic approach inspired by the ideas in [21]. Namely, we construct a graph \mathcal{H} whose vertices are agents, and there is an edge between two agents e and e' if removing both e and e' results in a system with no feasible solution. We call \mathcal{H} the dependency graph of the pruned system. We then compute the largest eigenvalue of \mathcal{H} (or, more precisely, the maximum of the largest eigenvalues of its connected components), which we denote by $\alpha_{\mathcal{H}}$, and scale the bid of each agent by the respective coordinate of the eigenvector that corresponds to $\alpha_{\mathcal{H}}$.

A given set system may be pruned in different ways, thus leading to different values of $\alpha_{\mathcal{H}}$. We will refer to the largest of them, i.e., $\alpha = \sup_{\mathcal{H}} \alpha_{\mathcal{H}}$, as the *eigenvalue* of our set system. It turns out that this quantity plays an important role in our analysis.

We show that the r -out-of- k -system mechanism and the $\sqrt{\cdot}$ -mechanism for the single path problem that were presented in [21] can be viewed as instantiations of our PRUNING-LIFTING MECHANISM. We then apply our scheme to two other classes of set systems: vertex cover systems, where the goal is to buy a vertex cover in a given graph, and k -path systems, where the goal is to buy k edge-disjoint paths between two vertices of a given graph.

The k -path problem generalizes both the r -out-of- k problem and the single path problem, and captures many other natural scenarios. However, this problem received limited attention from the algorithmic mechanism design community so far (see, however, [19]), perhaps due to its inherent difficulty: the interactions among the agents can be quite complex, and, prior to this work, it was not known how to characterize Nash equilibria of the first-price auctions for this setting in terms of the network structure. In this paper, we obtain a strong lower bound on the total payments in Nash equilibria. We then use this bound to show that a natural variant of the PRUNING-LIFTING MECHANISM that prunes all edges except those in the cheapest flow of size $k + 1$ has frugality ratio $\alpha \frac{k+1}{k}$. Moreover, we show that this bound can be improved by a factor of $k + 1$ if we consider a weaker payment bound suggested in [10], which corresponds to a buyer-pessimal rather than buyer-optimal Nash equilibrium (i.e., the difference between two frugality bounds is akin to that between the price of anarchy and the price of stability).

For the vertex cover problem, an earlier paper [10] described a mechanism with frugality ratio 2Δ , where Δ is the maximum degree of the input graph. Our approach results in a mechanism whose frugality ratio equals to the largest eigenvalue α of the adjacency matrix of the input graph. As $\alpha \leq \Delta$ for any graph G , this means that we improve the result of [10] by at least a factor of 2 for all graphs. Surprisingly, this stronger bound can be obtained by a simple modification of the analysis in [10].

Lower Bounds. We complement the bounds on the frugality of the PRUNING-LIFTING MECHANISM by proving strong lower bounds on the frugality of (almost) any truthful mechanism. In more detail,

we exhibit a family of cost vectors on which the payment of any *measurable* truthful mechanism can be lower-bounded in terms of α , where we call a mechanism measurable if the payment to any agent — as a function of other agents’ bids — is a Lebesgue measurable function. Lebesgue measurability is a much weaker condition than continuity or monotonicity; indeed, a mechanism that does not satisfy this condition is unlikely to be practically implementable! Our argument relies on Young’s inequality and applies to any set system.

To turn this lower bound on payments into a lower bound on frugality, we need to understand the structure of Nash equilibria for the bid vectors employed in our proof. For k -path systems, we can achieve this by using our characterization of Nash equilibria in such systems. As a result, we obtain a lower bound on frugality of any “measurable” truthful mechanism that shows that our mechanism is within a factor of $(k + 1)$ from optimal. Moreover, it is, in fact, optimal, with respect to the weaker payment bound of [10]. For r -out-of- k systems and single path systems, our bound improves the lower bounds on frugality given in [21] by a factor of 2 and $\sqrt{2}$, respectively. Our results give strong evidence that simply choosing the cheapest $(k + 1)$ -flow mechanism, which generalizes the $\sqrt{\cdot}$ -mechanism [21] for $k = 1$, is indeed an optimal frugal mechanism for paths systems.

For the vertex cover problem, characterizing the Nash equilibria turns out to be a more difficult task: in this case, the graph \mathcal{H} is equal to the input graph, and therefore is not guaranteed to have any regularity properties. However, we can still obtain non-trivial upper bounds on the payments in Nash equilibria. These bounds enable us to show that our mechanism for vertex cover is optimal for all triangle-free graphs, and, more generally, for all graphs that satisfy a simple local sparsity condition.

1.2 Related Work

There is a substantial literature on designing mechanisms with small payment for shortest path systems [2, 11, 12, 8, 9, 19, 29] as well as for other set systems [27, 6, 4, 21, 10], starting with the seminal work of Nisan and Ronen [26]. Our work is most closely related to [21], [10] and [29]: we employ the frugality benchmark defined in [21], improve the bounds of [21] and [10], and generalize the result of [29].

Simultaneously and independently, the idea of bounding frugality ratios of set system auctions in terms of eigenvalues of certain matrices was proposed by Kempe et al. [23]. In contrast with our work, in [23] the authors only study the frugality ratio of their mechanisms with respect to the relaxed payment bound (see Section 5). Their approach results in a 2-competitive mechanism for vertex cover systems, $2(k+1)$ -competitive mechanism for k -path systems, and a 4-competitive mechanism for cut auctions.

2 Preliminaries

A *set system* $(\mathcal{E}, \mathcal{F})$ is given by a set \mathcal{E} of *agents* and a collection $\mathcal{F} \subseteq 2^{\mathcal{E}}$ of *feasible sets*. We restrict our attention to *monopoly-free* set systems, i.e., we require $\bigcap_{S \in \mathcal{F}} S = \emptyset$. Each agent $e \in \mathcal{E}$ has a privately known *cost* $c(e)$ that represents the expenses that agent e incurs if he is involved in performing the task.

A *mechanism* for a set system $(\mathcal{E}, \mathcal{F})$ takes a *bid vector* $\mathbf{b} = (b(e))_{e \in \mathcal{E}}$ as input, where $b(e) \geq c(e)$ for any $e \in \mathcal{E}$, and outputs a set of *winners* $S \in \mathcal{F}$ and a *payment* $p(e)$ for each $e \in \mathcal{E}$. We require mechanisms to satisfy *voluntary participation*, i.e. $p(e) \geq b(e)$ for each $e \in S$ and $p(e) = 0$ for each $e \notin S$.

Given the output of a mechanism, the *utility* of an agent e is $p(e) - c(e)$ if e is a winner and 0 otherwise. We assume that agents are rational, i.e. aim to maximize their own utility. Thus, they may lie about their true costs, i.e. bid $b(e) \neq c(e)$ if they can profit by doing so. We say that a mechanism is *truthful* if every agent maximizes his utility by bidding his true value, no matter what bids other agents submit. A weaker solution concept is that of *Nash equilibrium*: a bid vector constitutes a (pure) Nash

equilibrium if no agent can increase his utility by unilaterally changing his bid. Nash equilibria describe stable states of the market and can be seen as natural outcomes of a bargaining process.

There is a well-known characterization of winner selection rules that yield truthful mechanisms.

Theorem 2.1 ([24, 2]). *A mechanism is truthful if and only if its winner selection rule is monotone, i.e., no losing agent can become a winner by increasing his bid, given the fixed bids of all other agents. Further, for a given monotone selection rule, there is a unique truthful mechanism with this selection rule: the payment to each winner is his threshold bid, i.e. the highest value he could bid and still win.*

An example of a truthful set system auction is given by the VCG mechanism [28, 7, 18]. However, as discussed in Section 1, VCG often results in a large overpayment to winners. Another natural mechanism for buying a set is the *first-price auction*: given the bid vector \mathbf{b} , pick a subset $S \in \mathcal{F}$ minimizing $b(S)$, and pay each winner $e \in S$ his bid $b(e)$. While the first-price auction is not truthful, and more generally, does not possess dominant strategies, it essentially admits a Nash equilibrium with a relatively small total payment. (More accurately, as observed by [19], a first-price auction may not have a pure strategy Nash equilibrium. However, this non-existence result can be circumvented in several ways, e.g. by considering instead an ε -Nash equilibrium for arbitrarily small $\varepsilon > 0$ or using oracle access to the true costs of agents to break ties.) The payment in a buyer-optimal Nash equilibrium would constitute a natural benchmark for truthful mechanisms. However, due to the difficulties described above, we use instead the following benchmark proposed by Karlin et al. [21], which captures the main properties of a Nash equilibrium.

Definition 2.2 (Benchmark $\nu(\mathbf{c})$ [21]). *Given a set system $(\mathcal{E}, \mathcal{F})$, and a feasible set $S \in \mathcal{F}$ of minimum total cost w.r.t. \mathbf{c} , let $\nu(\mathbf{c})$ be the value of an optimal solution to the following optimization problem:*

$$\begin{aligned} \min \quad & \sum_{e \in S} b(e) \\ \text{s.t.} \quad & (1) \ b(e) \geq c(e) \text{ for all } e \in \mathcal{E} \\ & (2) \ \sum_{e \in S \setminus T} b(e) \leq \sum_{e \in T \setminus S} c(e) \text{ for all } T \in \mathcal{F} \\ & (3) \ \text{For every } e \in S \text{ there is a } T \in \mathcal{F} \text{ s.t. } e \notin T \text{ and } \sum_{e' \in S \setminus T} b(e') = \sum_{e' \in T \setminus S} c(e') \end{aligned}$$

Intuitively, in the optimal solution of the above system, S is the set of winners in the first-price auction. By condition (3), no winner $e \in S$ can improve his utility by increasing his bid $b(e)$, as he would not be a winner anymore. In addition, by conditions (1) and (2), no agent $e \in \mathcal{E} \setminus S$ can obtain a positive utility by decreasing his bid. Hence, $\nu(\mathbf{c})$ gives the value of the cheapest Nash equilibrium of the first-price auction assuming that the most “efficient” feasible set S wins.

Definition 2.3 (Frugality Ratio). *Let \mathcal{M} be a truthful mechanism for the set system $(\mathcal{E}, \mathcal{F})$ and let $p_{\mathcal{M}}(\mathbf{c})$ denote the total payment of \mathcal{M} when the true costs are given by a vector \mathbf{c} . Then the frugality ratio of \mathcal{M} on \mathbf{c} is defined as $\phi_{\mathcal{M}}(\mathbf{c}) = \frac{p_{\mathcal{M}}(\mathbf{c})}{\nu(\mathbf{c})}$. Further, the frugality ratio of \mathcal{M} is defined as $\phi_{\mathcal{M}} = \sup_{\mathbf{c}} \phi_{\mathcal{M}}(\mathbf{c})$.*

3 Pruning-Lifting Mechanism

In this section, we describe a general scheme for designing truthful mechanisms for set systems, which we call PRUNING-LIFTING MECHANISM. For a given set system $(\mathcal{E}, \mathcal{F})$, the mechanism is composed of the following steps:

- **Pruning.** The goal of the pruning process is to drop some elements of \mathcal{E} to expose the structure of the competition between the agents; we denote the set of surviving agents by \mathcal{E}^* . We require the process to satisfy the following properties:
 - **Monotonicity:** for any given vector of other agents' bids, if an agent e is dropped when he bids b , he is also dropped if he bids any $b' > b$. We set $t_1(e) = \inf\{b' \mid e \text{ is dropped when bidding } b'\}$.
 - **Bid-independence:** for any given vector of other agents' bids, let b and b' be two bids of agent e such that e is not dropped when he submits either of them. Then for both of these bids the set \mathcal{E}^* of remaining agents is the same. That is, e cannot control the outcome of the pruning process as long as he survives. Monotonicity and bid-independence conditions are important to ensure the truthfulness of the mechanism.
 - **Monopoly-freeness:** the remaining set system must remain monopoly-free, i.e., $\bigcap_{S \in \mathcal{F}^*} S = \emptyset$, where $\mathcal{F}^* = \{S' \in \mathcal{F} \mid S' \subseteq \mathcal{E}^*\}$. This condition is necessary because in the winner selection stage we will choose a winning feasible set from \mathcal{F}^* . Therefore, we have to make sure that no winning agent can charge an arbitrarily high price due to lack of competition.
- **Lifting.** The goal of the lifting process is to assign a weight to each agent in \mathcal{E}^* in order to take into account the size of each feasible set. To this end, construct an undirected graph \mathcal{H} by (a) introducing a node v_e for each $e \in \mathcal{E}^*$, and (b) connecting v_e and $v_{e'}$ if and only if any feasible set in \mathcal{F}^* contains either e or e' . We will refer to \mathcal{H} as the *dependency graph* of \mathcal{E}^* . For each connected component \mathcal{H}_j of \mathcal{H} , compute the largest eigenvalue α_j of its adjacency matrix A_j , and let $(w(v_e))_{v_e \in \mathcal{H}_j}$ be the eigenvector of A_j associated with α_j . That is, $A_j \mathbf{w}^j = \alpha_j \mathbf{w}^j$, where $\mathbf{w}^j = ((w(v_e))_{v_e \in \mathcal{H}_j})^T$. Set $\alpha = \max \alpha_j$.
- **Winner selection.** Define $b'(e) = \frac{b(e)}{w(v_e)}$ for each $e \in \mathcal{E}^*$, and select a feasible set $S \in \mathcal{F}^*$ with the smallest total bids w.r.t. \mathbf{b}' . Let $t_2(e)$ be the threshold bid for $e \in \mathcal{E}^*$ to be selected at this stage.
- **Payment.** The payment to each winner $e \in S$ is $p(e) = \min\{t_1(e), t_2(e)\}$, where $t_1(e)$ and $t_2(e)$ are the two thresholds defined above.

Recall that the largest eigenvalue of the adjacency matrix of a connected graph is positive and its associated eigenvector has strictly positive coordinates [15]. Therefore, $w(v_e) > 0$ for all $e \in \mathcal{E}^*$.

We will now define a quantity $\alpha_{(\mathcal{E}, \mathcal{F})}$ that will be instrumental in characterizing the frugality ratio of truthful mechanisms on $(\mathcal{E}, \mathcal{F})$. Let $\mathcal{S}(\mathcal{E}, \mathcal{F})$ be the collection of all monopoly-free subsets of \mathcal{E} , i.e., set $\mathcal{S}(\mathcal{E}, \mathcal{F}) = \{S \subseteq \mathcal{E} \mid \bigcap_{T \in \mathcal{F}, T \subseteq S} T = \emptyset\}$. The elements of $\mathcal{S}(\mathcal{E}, \mathcal{F})$ are the possible outcomes of the pruning stage. For any subset $S \in \mathcal{S}(\mathcal{E}, \mathcal{F})$, let \mathcal{H}_S be its dependency graph and A_S be the adjacency matrix of \mathcal{H}_S . Let α_S be the largest eigenvalue of A_S (or the maximum of the largest eigenvalues of the adjacency matrices of the connected components of \mathcal{H}_S , if \mathcal{H}_S is not connected). Set $\alpha_{(\mathcal{E}, \mathcal{F})} = \max_{S \in \mathcal{S}(\mathcal{E}, \mathcal{F})} \alpha_S$; we will refer to $\alpha_{(\mathcal{E}, \mathcal{F})}$ as the *eigenvalue of the set system* $(\mathcal{E}, \mathcal{F})$.

Note that once $\mathcal{E}^* \in \mathcal{S}(\mathcal{E}, \mathcal{F})$ is selected in the pruning step, the computation of α and the weight vector $(w(v_e))_{e \in \mathcal{E}^*}$ does not depend on the bid vector. This property is crucial for showing that our mechanism is truthful. (Due to space limits, most proofs in the paper are relegated to the Appendix.)

Theorem 3.1. PRUNING-LIFTING MECHANISM *is truthful for any set system* $(\mathcal{E}, \mathcal{F})$.

In the rest of this section, we will show that the mechanisms for r -out-of- k systems and single path systems proposed in [21] can be viewed as instantiations of our PRUNING-LIFTING MECHANISM. By Theorems 2.1 and 3.1, we can ignore the payment rule in the following discussion.

3.1 r -out-of- k Systems Revisited

In an r -out-of- k system, the set of agents \mathcal{E} is a union of k disjoint subsets S_1, \dots, S_k and the feasible sets are unions of exactly r of those subsets. Given a bid vector \mathbf{b} , renumber the subsets S_1, \dots, S_k in order of non-decreasing bids, i.e., $b(S_1) \leq b(S_2) \leq \dots \leq b(S_k)$.

The mechanism proposed in [21] deletes all but the first $r + 1$ subsets, and then solves a system of equations given by

$$(\diamond) \quad \beta = \frac{1}{rx_i} \cdot \sum_{j \neq i} x_j \cdot |S_j| \quad \text{for } i = 1, \dots, r + 1.$$

It then scales the bid of each set S_i by setting $b'(S_i) = \frac{b(S_i)}{x_i}$, discards the set with the highest scaled bid w.r.t. \mathbf{b}' , and outputs the remaining sets.

Now, clearly, the first step of this mechanism can be interpreted as a pruning stage. Further, for r -out-of- k systems the graph \mathcal{H} constructed in the lifting stage of our mechanism is a complete $(r + 1)$ -partite graph. It is not hard to verify that for any positive solution $(x_1, \dots, x_{r+1}, \beta)$ of the system (\diamond) , $\beta \cdot r$ gives the largest eigenvalue of the adjacency matrix of \mathcal{H} and $(x_1, \dots, x_1, \dots, x_{r+1}, \dots, x_{r+1})$ is the corresponding eigenvector. Thus, the mechanism of [21] implements PRUNING-LIFTING MECHANISM for r -out-of- k systems.

In [21] it is shown that the frugality ratio of this mechanism is β , and that the frugality ratio of any truthful mechanism for r -out-of- k systems is at least $\frac{\beta}{2}$. As r -out-of- k systems can be viewed as a special case of r -path systems, Theorem 6.5 allows us to improve this lower bound to $\frac{\beta r}{r} = \beta$.

3.2 Single Path Mechanisms Revisited

In a single path system, agents are edges of a given directed graph $G = (V, E)$ with two specified vertices s and t , i.e. $\mathcal{E} = E$ and \mathcal{F} consists of all sets of edges that contain a path from s to t .

Given a bid vector \mathbf{b} , the $\sqrt{\cdot}$ -mechanism [21] first selects two edge-disjoint s - t paths P and P' that minimize $b(P) + b(P')$. Assume that P and P' intersect at $s = v_1, v_2, \dots, v_{\ell+1} = t$, in the order in which they appear in P and P' . Let P_i and P'_i be the subpaths of P and P' from v_i to v_{i+1} , respectively. The $\sqrt{\cdot}$ -mechanism sets $b'(e) = b(e)\sqrt{|P_i|}$ for $e \in P_i$, $b'(e) = b(e)\sqrt{|P'_i|}$ for $e \in P'_i$, and chooses a cheapest path in $P \cup P'$ w.r.t. \mathbf{b}' .

As in the previous case, the selection of P and P' can be viewed as the pruning process. The corresponding graph \mathcal{H} consists of ℓ connected components, where the i -th component \mathcal{H}_i is a complete bipartite graph with parts of size $|P_i|$ and $|P'_i|$. Its largest eigenvalue is given by $\alpha_i = \sqrt{|P_i||P'_i|}$, and the coordinates of the corresponding eigenvector are given by $w(v_e) = 1/\sqrt{|P_i|}$ for $e \in P_i$ and $w(v_e) = 1/\sqrt{|P'_i|}$ for $e \in P'_i$. Thus, the $\sqrt{\cdot}$ -mechanism can be viewed as a special case of the PRUNING-LIFTING MECHANISM. It is shown that the frugality ratio of the $\sqrt{\cdot}$ -mechanism is within a factor of $2\sqrt{2}$ from optimal; Theorem 6.5 below shows that this bound can be improved by a factor of $\sqrt{2}$ (this has also been shown by Yan [29] via a proof that is considerably more complicated than ours).

4 Vertex Cover Systems

In the vertex cover problem, we are given a graph $G = (V, E)$ whose vertices are owned by selfish agents. Our goal is to purchase a vertex cover of G . That is, we have $\mathcal{E} = V$, and \mathcal{F} is the collection of all vertex covers of G . Let A denote the adjacency matrix of G , and let Δ , $\alpha = \alpha_{(\mathcal{E}, \mathcal{F})}$ and $\mathbf{w} = (w(v))_{v \in V}$ denote, respectively, the maximum degree of G , the largest eigenvalue of A and the corresponding eigenvector.

We will use the pruning-lifting scheme to construct a mechanism whose frugality ratio is α ; this improves the bound of 2Δ given in [10] by at least a factor of 2 for all graphs, and by as much as a factor of $\Theta(\sqrt{n})$ for some graphs (e.g., the star).

Observe first that the vertex cover system plays a special role in the analysis of the performance of the pruning-lifting scheme. Indeed, on one hand, it is straightforward to apply the PRUNING-LIFTING MECHANISM to this system: since removing any agent will make each of its neighbors a monopolist, the pruning stage of our scheme is redundant, i.e., $\mathcal{H} = G$. That is, there is a unique implementation of PRUNING-LIFTING MECHANISM for vertex cover systems: we set $b'(v) = \frac{b(v)}{w(v)}$ for all $v \in V$, pick any $S \in \arg \min\{b'(T) \mid T \text{ is a vertex cover for } G\}$ to be the winning set, and pay each agent $v \in S$ his threshold bid $t(v)$. On the other hand, for general set systems, any feasible set in the pruned system corresponds to a vertex cover of \mathcal{H} : indeed, by construction of the graph \mathcal{H} , any feasible set must contain at least one endpoint of any edge of \mathcal{H} . In general, the converse is not true: a vertex cover of \mathcal{H} is not necessarily a feasible set. However, for k -path systems it is possible to show that any cover of \mathcal{H} corresponds to a k -flow. We can use this fact to provide an alternative proof of Theorems 5.1 and 5.6.

We will now bound the frugality of PRUNING-LIFTING MECHANISM for vertex cover systems.

Theorem 4.1. *The frugality ratio of PRUNING-LIFTING MECHANISM for vertex cover systems on a graph G is at most $\alpha = \alpha_{(\mathcal{E}, \mathcal{F})}$.*

Proof. Since our mechanism is truthful, we have $b(v) = c(v)$ for all $v \in V$. By optimality of S we have $b'(v) \leq \sum_{uv \in E, u \notin S} b'(u)$, and therefore $t(v) \leq w(v) \sum_{uv \in E, u \notin S} \frac{c(u)}{w(u)}$. Thus, we can bound the total payment of our mechanism given a bid vector \mathbf{b} as

$$\sum_{v \in S} t(v) \leq \sum_{v \in S} w(v) \sum_{uv \in E, u \notin S} \frac{c(u)}{w(u)} = \sum_{u \notin S} \frac{c(u)}{w(u)} \sum_{uv \in E} w(v) = \sum_{u \notin S} \frac{c(u)}{w(u)} \alpha w(u) = \alpha \sum_{u \notin S} c(u).$$

Lemma 8 in [10] shows that for any cost vector \mathbf{c} we have $\nu(\mathbf{c}) \geq \sum_{u \notin S} c(u)$. Therefore, the frugality ratio of PRUNING-LIFTING MECHANISM for vertex cover on G is at most α . \square

In Section 6.1 we show that our mechanism is optimal for a large class of graphs.

4.1 Computational Issues

To implement PRUNING-LIFTING MECHANISM for vertex cover, we need to select the vertex cover that minimizes the scaled costs given by $(b'(v))_{v \in V}$, i.e., to solve an NP-hard problem. However, the argument in Theorem 4.1 applies to any truthful mechanism that selects a *locally optimal* solution, i.e., a vertex cover S that satisfies $b'(v) \leq \sum_{uv \in E, u \notin S} b'(u)$ for all $v \in S$. Paper [10] argues that any monotone winner selection algorithm for vertex cover can be transformed into a locally optimal one, and shows that a variant of the classic 2-approximation algorithm for this problem [3] is monotone. This leads to the following corollary.

Corollary 4.2. *There exists a truthful polynomial-time vertex cover auction that given a graph G outputs a solution whose cost is within a factor of 2 from optimal and whose frugality ratio is at most α .*

5 Multiple Paths Systems

In this section, we study in detail k -paths systems for a given integer $k \geq 1$. In these systems, the set of agents \mathcal{E} is the set of edges of a directed graph $G = (V, E)$ with two specified vertices $s, t \in V$. The

feasible sets are sets of edges that contain k edge-disjoint s - t paths. Clearly, these set systems generalize both r -out-of- k systems and single path systems.

Our mechanism for k -paths systems for a given directed graph G , which we call PRUNING-LIFTING k -PATHS MECHANISM, is a natural generalization of the $\sqrt{\cdot}$ -mechanism [21]: In the pruning stage of our mechanism, given a bid vector \mathbf{b} , we pick $k+1$ edge-disjoint s - t paths P_1, \dots, P_{k+1} so as to minimize their total bid w.r.t. the bid vector \mathbf{b} . Clearly, this procedure is monotone and bid-independent. Let $G^*(\mathbf{b})$ denote the subgraph composed of these $k+1$ paths. The remaining steps of the mechanism (lifting, winner selection, payment determination) are the same as in the general case (Section 3). Since the PRUNING-LIFTING k -PATHS MECHANISM is an implementation of the PRUNING-LIFTING MECHANISM, Theorem 3.1 implies that it is truthful.

Let \mathcal{G}^{k+1} denote the set of all subgraphs of G that can be represented as a union of $k+1$ edge-disjoint s - t paths in G . For any $G^* \in \mathcal{G}^{k+1}$, let $\mathcal{H}(G^*)$ denote the dependency graph of G^* , and let $\alpha(G^*)$ denote the maximum of the largest eigenvalues of the connected components of $\mathcal{H}(G^*)$. Set $\alpha_{k+1} = \max\{\alpha(G^*) \mid G^* \in \mathcal{G}^{k+1}\}$. We can bound the frugality ratio of our mechanism as follows.

Theorem 5.1. *The frugality ratio of PRUNING-LIFTING k -PATHS MECHANISM is at most $\alpha_{k+1} \frac{k+1}{k}$.*

To prove Theorem 5.1, we need the following definition.

Definition 5.2 (Minimum Longest Path $\delta_{k+1}(G, \mathbf{c})$). *For any $k+1$ edge-disjoint s - t paths P_1, \dots, P_{k+1} in a directed graph G , let $\delta_{k+1}(P_1, \dots, P_{k+1}, \mathbf{c})$ denote the length of the longest s - t path w.r.t. cost vector \mathbf{c} in the subgraph G' composed of P_1, \dots, P_{k+1} (if G' contains a positive length cycle, set $\delta_{k+1}(P_1, \dots, P_{k+1}) = +\infty$). Define*

$$\delta_{k+1}(G, \mathbf{c}) = \min\{\delta_{k+1}(P_1, \dots, P_{k+1}, \mathbf{c}) \mid P_1, \dots, P_{k+1} \text{ are } k+1 \text{ edge-disjoint } s\text{-}t \text{ paths}\}.$$

It turns out that we can give a lower bound on $\nu(\mathbf{c})$ in terms of $\delta_{k+1}(G, \mathbf{c})$.

Lemma 5.3. *For any k -paths system on a given graph G with costs \mathbf{c} , we have $\nu(\mathbf{c}) \geq k \cdot \delta_{k+1}(G, \mathbf{c})$.*

Let $L(G, \mathbf{c})$ be the length of the longest path in $G^*(\mathbf{c})$, where $G^*(\mathbf{c})$ is the output of our pruning process on the bid vector \mathbf{c} . Our second lemma gives an upper bound on the payment of our mechanism in terms of $L(G, \mathbf{c})$.

Lemma 5.4. *For any k -paths system on a given graph G with costs \mathbf{c} , the total payment of PRUNING-LIFTING k -PATHS MECHANISM on \mathbf{c} is at most $\alpha(G^*(\mathbf{c}))L(G, \mathbf{c})$.*

Theorem 5.1 now follows from Lemmas 5.3 and 5.4 and the observation that $G^*(\mathbf{c}) \in \mathcal{G}^{k+1}$ and $L(G, \mathbf{c}) \leq (k+1)\delta_{k+1}(G, \mathbf{c})$ (see Appendix B for details).

In Section 6.2, we show a lower bound of $\frac{\alpha_{k+1}}{k}$ on the frugality ratio of essentially any truthful k -path mechanism; thus the frugality ratio of our mechanism is within a factor of $k+1$ from optimal. This gap leaves the question of whether there is a better truthful mechanism for k -paths systems. One might hope that a different pruning approach could lead to a smaller frugality ratio. In particular, the proof of Theorem 5.1 suggests that we could get a stronger result by pruning the graph so as to minimize the length of the longest path $\delta_{k+1}(G^*, \mathbf{c})$ in the surviving graph G^* . While the argument above shows that — under truthful bidding — such mechanism would have an optimal frugality ratio, unfortunately, it turns out that this pruning process is not monotone [22].

We can, however, show that our mechanism is *optimal* with respect to a weaker benchmark, namely, one that corresponds to a buyer-pessimal rather than buyer-optimal Nash equilibrium. This benchmark was introduced in [10], and has been recently used by Kempe et al. [23]. As argued in [10] and [23], unlike ν , this benchmark enjoys natural monotonicity properties and is easier to work with.

Definition 5.5 (Benchmark $\mu(\mathbf{c})$ [10]). *Given a set system $(\mathcal{E}, \mathcal{F})$, and a feasible set $S \in \mathcal{F}$ of minimum total cost w.r.t. \mathbf{c} , let $\mu(\mathbf{c})$ be the value of an optimal solution to the following optimization problem:*

$$\begin{aligned}
& \max && \sum_{e \in S} b(e) \\
& \text{s.t.} && (1) \ b(e) \geq c(e) \text{ for all } e \in \mathcal{E} \\
& && (2) \ \sum_{e \in S \setminus T} b(e) \leq \sum_{e \in T \setminus S} c(e) \text{ for all } T \in \mathcal{F} \\
& && (3) \ \text{For every } e \in S \text{ there is a } T \in \mathcal{F} \text{ s.t. } e \notin T \text{ and } \sum_{e' \in S \setminus T} b(e') = \sum_{e' \in T \setminus S} c(e')
\end{aligned}$$

We will refer to the quantity $\sup_{\mathbf{c}} \frac{p_{\mathcal{M}}(\mathbf{c})}{\mu(\mathbf{c})}$ as the μ -frugality ratio of a mechanism \mathcal{M} , where $p_{\mathcal{M}}(\mathbf{c})$ is the total payment of mechanism \mathcal{M} on a bid vector \mathbf{c} .

The programs for $\nu(\mathbf{c})$ and $\mu(\mathbf{c})$ differ in their objective function only: while $\nu(\mathbf{c})$ minimizes the total payment, $\mu(\mathbf{c})$ maximizes it. In particular, this means that in the program for $\mu(\mathbf{c})$ we can omit constraint (3), i.e., $\mu(\mathbf{c})$ can be obtained as a solution to a *linear* program. Kempe et al. [23] show that the μ -frugality ratio of PRUNING-LIFTING k -PATHS MECHANISM is within a factor of $2(k+1)$ from optimal. Our next result, combined with the observation that our lower bound on the performance of “all” truthful mechanisms also holds for the μ -frugality ratio, shows that PRUNING-LIFTING k -PATHS MECHANISM is, in fact, optimal with respect to μ . The proof proceeds by constructing a bid vector \mathbf{y} that satisfies constraints (1) and (2) in the definition of $\mu(\mathbf{c})$ and pays at least $kL(G, \mathbf{c})$, and uses the observation that for any fixed network the cost of a flow of size x is a convex piecewise-linear function of x . Details of the proof are given in Appendix C.

Theorem 5.6. *The μ -frugality ratio of PRUNING-LIFTING k -PATHS MECHANISM is at most $\frac{\alpha_{k+1}}{k}$.*

6 Lower Bounds

We say that a mechanism \mathcal{M} for a set system $(\mathcal{E}, \mathcal{F})$ is *measurable* if the payment $p(e)$ of any agent $e \in \mathcal{E}$ is a Lebesgue measurable function of all agents’ bids. We will now use Young’s inequality to give a lower bound on total payments of any measurable truthful mechanism with bounded frugality ratio.

Theorem 6.1 (Young’s inequality). *Let $f_1 : [0, a] \rightarrow \mathbb{R}^+ \cup \{0\}$, $f_2 : [0, b] \rightarrow \mathbb{R}^+ \cup \{0\}$ be two Lebesgue measurable functions that are bounded on their domain. Assume that whenever $y > f_1(x)$ for some $0 < x \leq a$, $0 < y \leq b$, we have $x \leq f_2(y)$. Then*

$$\int_0^a f_1(x) dx + \int_0^b f_2(y) dy \geq ab.$$

Fix a set system $(\mathcal{E}, \mathcal{F})$ with $|\mathcal{E}| = n$ and let $S_{(\mathcal{E}, \mathcal{F})} \in \mathcal{S}(\mathcal{E}, \mathcal{F})$ be a subset with $\alpha_S = \alpha_{(\mathcal{E}, \mathcal{F})}$. For any $e \in S_{(\mathcal{E}, \mathcal{F})}$, let $\mathbf{c}_{e,x}$ denote a bid vector where e bids x , all agents in $S_{(\mathcal{E}, \mathcal{F})} \setminus \{e\}$ bid 0, and all agents in $\mathcal{E} \setminus S_{(\mathcal{E}, \mathcal{F})}$ bid $n+1$.

Lemma 6.2. *For any set system $(\mathcal{E}, \mathcal{F})$ and any measurable truthful mechanism \mathcal{M} with bounded frugality ratio, there exists an agent $e \in S_{(\mathcal{E}, \mathcal{F})}$ and a real value $0 < x \leq 1$ such that the total payment of \mathcal{M} on bid vector $\mathbf{c}_{e,x}$ is at least $\alpha_{(\mathcal{E}, \mathcal{F})}x$.*

Proof. Set $S = S_{(\mathcal{E}, \mathcal{F})}$, $\mathcal{H} = \mathcal{H}_S$, $A = A_S$, $\alpha = \alpha_S = \alpha_{(\mathcal{E}, \mathcal{F})}$. We will assume from now on that $\mathcal{H} = (S, E(\mathcal{H}))$ is connected; if this is not the case, our argument can be applied without change to the connected component of \mathcal{H} that corresponds to α . Let $\mathbf{w} = (w_v)_{v \in S}$ be the eigenvector of A that is associated with α . By normalization, we can assume that $\max_{v \in S} w_v = 1$.

The proof is by contradiction: assume that there is a truthful mechanism \mathcal{M} that pays less than αx on any bid vector of the form $\mathbf{c}_{e,x}$ for all $e \in S$ and all $0 < x \leq 1$. Recall that for any such bid vector the cost of each agent in $\mathcal{E} \setminus S$ is $n + 1$. Since $\alpha \leq n$ and $x \leq 1$, this implies that \mathcal{M} never picks any agents from $\mathcal{E} \setminus S$ on any $\mathbf{c}_{e,x}$, i.e., effectively \mathcal{M} operates on S . For any edge vu of \mathcal{H} and any $x > 0$, let $p_{uv}(x)$ denote the payment to v on the bid vector $\mathbf{c}_{u,x}$. Observe that measurability of \mathcal{M} implies that $p_{uv}(x)$ is measurable (since it is a restriction of a measurable function). In this notation, our assumption can be restated as

$$\sum_{uv \in E(\mathcal{H})} p_{uv}(x) < \alpha x \quad (1)$$

for all $u \in S$ and any $0 < x \leq 1$.

It is easy to see that given a bid vector $\mathbf{c}_{u,z}$ with $z \leq 1$, \mathcal{M} never selects u as a winner. Indeed, suppose that u wins given $\mathbf{c}_{u,z}$. Then by the truthfulness of \mathcal{M} , if we reduce u 's true cost from z to 0, u still wins and receives a payment of at least z . Since the set system restricted to S is monopoly-free, the resulting cost vector \mathbf{c}' satisfies conditions (1)–(3) in the definition of ν , and hence $\nu(\mathbf{c}') = 0$. Thus the frugality ratio of \mathcal{M} is $+\infty$, a contradiction. By the construction of \mathcal{H} , this means that any $v \in S$ with $uv \in E(\mathcal{H})$ wins given $\mathbf{c}_{u,z}$.

Now, fix some x, y such that $0 < x, y \leq 1$ and $y > p_{vu}(x)$, and consider a situation where v bids x , u bids y , all agents in $S \setminus \{u, v\}$ bid 0, and all agents in $\mathcal{E} \setminus S$ bid $n + 1$. Clearly, in this situation agent u loses and thus v wins with a payment of at least x . By the truthfulness of \mathcal{M} , the same holds if v lowers his bid to 0. Thus, for any $0 < x, y \leq 1$, $y > p_{vu}(x)$ implies $p_{uv}(y) \geq x$.

By our assumption, we have $p_{uv}(x) \leq \alpha x$, $p_{vu}(x) \leq \alpha x$ for $x \in [0, 1]$. Hence, for any $uv \in E(\mathcal{H})$ the functions $p_{uv}(x)$ and $p_{vu}(x)$ satisfy all conditions of Young's inequality on $[0, 1]$.

Let $A = (a_{uv})_{u,v \in S}$, and consider the scalar product $\langle \mathbf{w}, A\mathbf{w} \rangle = \langle \mathbf{w}, \alpha \mathbf{w} \rangle = \alpha \langle \mathbf{w}, \mathbf{w} \rangle$. We have $\langle \mathbf{w}, A\mathbf{w} \rangle = \sum_{uv \in E(\mathcal{H})} w_u w_v$. As we normalized \mathbf{w} so that $w_u, w_v \leq 1$, by Young's inequality, we can bound $w_u w_v$ by $\int_0^{w_u} p_{uv}(x) dx + \int_0^{w_v} p_{vu}(x) dx$. Therefore,

$$\begin{aligned} \alpha \langle \mathbf{w}, \mathbf{w} \rangle &= \sum_{u,v \in S} a_{uv} w_u w_v \leq \sum_{u,v \in S} \left(\int_0^{w_u} a_{uv} p_{uv}(x) dx + \int_0^{w_v} a_{uv} p_{vu}(y) dy \right) \\ &= 2 \sum_{u \in S} \int_0^{w_u} \sum_{v \in S} a_{uv} p_{uv}(x) dx < 2\alpha \sum_{u \in S} \int_0^{w_u} x dx = \alpha \sum_{u \in S} w_u^2 = \alpha \langle \mathbf{w}, \mathbf{w} \rangle, \end{aligned}$$

where the last inequality follows from (1). This is a contradiction, so the proof is complete. \square

6.1 Vertex Cover Systems

For vertex cover systems, deleting any of the agents would result in a monopoly. Therefore, Lemma 6.2 simply says that for any measurable truthful mechanism \mathcal{M} on a graph $G = (V, E)$, there exists a $v \in V$ such that the total payment on bid vector $x \cdot \mathbf{c}_v$ is at least αx , where α is the largest eigenvalue of the adjacency matrix of G and \mathbf{c}_v is the cost vector given by $c_v(u) = 1$ if $u = v$, and $c_v(u) = 0$ if $u \in V \setminus \{v\}$.

Given a graph $G = (V, E)$ and a vertex $v \in V$, let L_v denote the set of all maximal cliques in G that contain v . Let ρ_v denote the size of the smallest clique in L_v .

Lemma 6.3. *We have $\nu(x \cdot \mathbf{c}_v) \leq x(\rho_v - 1)$ for any $x > 0$.*

Combining Lemma 6.3 with Lemma 6.2 yields the following result.

Theorem 6.4. *For any graph G , the frugality ratio of any measurable truthful vertex cover auction on G is at least $\frac{\alpha}{\rho}$, where α is the largest eigenvalue of the adjacency matrix of G , and $\rho = \max_{v \in V} \rho_v$.*

The bound given in Lemma 6.3 is not necessarily optimal; we can construct a family of graphs where for some vertex v the quantity ρ_v is linear in the size of the graph, while $\nu(\mathbf{c}_v) = O(1)$. Nevertheless, Theorem 6.4 shows that the mechanism described in Section 4 has optimal frugality ratio for, e.g., all triangle-free graphs and, more generally, all graphs G such that for each vertex $v \in G$, the subgraph induced on the neighbors of v contains an isolated vertex.

6.2 Multiple Path Systems

Let $(\mathcal{E}, \mathcal{F})$ be a k -path system on a graph $G = (V, E)$. Consider a set $S \in \mathcal{S}(\mathcal{E}, \mathcal{F})$ with $\alpha_S = \alpha_{(\mathcal{E}, \mathcal{F})}$. It is not hard to see that S is a union of $k + 1$ edge-disjoint paths; this follows, e.g., from the proof of Theorem B.1. Hence, we have $\alpha_{(\mathcal{E}, \mathcal{F})} = \alpha_{k+1}$.

As before, for any $e \in S$, let $\mathbf{c}_{e,x}$ denote the cost vector with $\mathbf{c}_{e,x}(e) = x$, $\mathbf{c}_{e,x}(u) = 0$ for all $u \in S \setminus \{e\}$, $\mathbf{c}_{e,x}(w) = n + 1$ for all $w \in E \setminus S$. It is easy to see that we have $\mu(\mathbf{c}_{e,x}) = \nu(\mathbf{c}_{e,x}) = kx$ for any $e \in S$, $x > 0$. Combining this observation with Lemma 6.2, we obtain the following result.

Theorem 6.5. *For any graph $G = (V, E)$, both the frugality ratio and the μ -frugality ratio of any measurable truthful k -path auction on G are at least $\frac{\alpha_{k+1}}{k}$.*

In Section 5, we show that the frugality ratio and the μ -frugality ratio of PRUNING-LIFTING k -PATHS MECHANISM are bounded by, respectively, $\alpha_{k+1} \frac{k+1}{k}$ and $\frac{\alpha_{k+1}}{k}$. Together with Theorem 6.5 this implies that PRUNING-LIFTING k -PATHS MECHANISM has optimal μ -frugality ratio, and its frugality ratio is within a factor of $(k + 1)$ from optimal.

7 Conclusions and Open Problems

In this paper, we propose a uniform scheme for designing frugal truthful mechanisms. We show that several existing mechanisms can be viewed as instantiations of our scheme, and describe its applications to k -path systems and vertex cover systems. We demonstrate that our scheme produces mechanisms with good frugality ratios for k -path systems and a large subclass of vertex cover systems; for k -path systems, we show that our mechanism has the optimal frugality ratio. Moreover, all mechanisms described in this paper are polynomial-time computable. We believe that our scheme can be applied to many other set systems, resulting in mechanisms with near-optimal frugality ratios.

It would be interesting to understand the limits of applicability of our scheme. Indeed, for some set systems the minimal monopoly-free subsystem does not necessarily exhibit a lot of connections between agents, i.e., the corresponding dependency graph is rather sparse. It seems that for such cases our scheme does not produce mechanisms with good frugality ratio. Formalizing this intuition and developing alternative approaches for designing frugal mechanisms in such settings is an interesting research direction.

8 Acknowledgements

We thank David Kempe, Christos Papadimitriou, and Yaron Singer for helpful discussions.

References

- [1] A. Archer, É. Tardos, *Truthful Mechanisms for One-Parameter Agents*, FOCS 2001, 482-491.
- [2] A. Archer, É. Tardos, *Frugal Path Mechanisms*, SODA 2002, 991-999.
- [3] R. Bar-Yehuda, S. Even, *A Local Ratio Theorem for Approximating the Weighted Vertex Cover Problem*, Annals of Discrete Mathematics, V.25, 27-46, 1985.
- [4] S. Bikhchandani, S. de Vries, J. Schummer, R. Vohra, *Linear Programming and Vickrey Auctions*, IMA Volume in Mathematics and its Applications, Mathematics of the Internet: E-auction and Markets, V.127, 75-116, 2001.
- [5] N. Chen, A. R. Karlin, *Cheap Labor Can Be Expensive*, SODA 2007, 707-715.
- [6] G. Calinescu, *Bounding the Payment of Approximate Truthful Mechanisms*, ISAAC 2004, 221-233.
- [7] E. H. Clarke, *Multipart Pricing of Public Goods*, Public Choice, V.11, 17-33, 1971.
- [8] A. Czumaj, A. Ronen, *On the Expected Payment of Mechanisms for Task Allocation*, PODC 2004, 98-106.
- [9] E. Elkind, *True Costs of Cheap Labor are Hard to Measure: Edge Deletion and VCG Payments in Graphs*, EC 2005, 108-116.
- [10] E. Elkind, L. A. Goldberg, P. W. Goldberg, *Frugality Ratios and Improved Truthful Mechanisms for Vertex Cover*, EC 2007, 336-345.
- [11] E. Elkind, A. Sahai, K. Steiglitz, *Frugality in Path Auctions*, SODA 2004, 701-709.
- [12] J. Feigenbaum, C. H. Papadimitriou, R. Sami, S. Shenker, *A BGP-Based Mechanism for Lowest-Cost Routing*, PODC 2002, 173-182.
- [13] L. R. Ford, D. R. Fulkerson, *Maximal Flow through a Network*, Canadian Journal of Mathematics, V.8, 399-404, 1956.
- [14] S. Fortune, J. Hopcroft, J. Wyllie, *The Directed Subgraph Homeomorphism Problem*, Theoretical Computer Science, V.10(2), 111-121, 1980.
- [15] C. Godsil, G. Royle, *Algebraic Graph Theory*, Graduate Texts in Mathematics, V.207, Springer, 2001.
- [16] A. V. Goldberg, R. E. Tarjan, *Finding Minimum Cost Circulation by Canceling Negative Cycles*, Journal of the ACM, V.36(4), 873-886, 1989.
- [17] N. Gravin, N. Chen, *A Note on k -Shortest Paths Problem*, accepted by Journal of Graph Theory.
- [18] T. Groves, *Incentives in Teams*, Econometrica, V.41, 617-631, 1973.
- [19] N. Immorlica, D. Karger, E. Nikolova, R. Sami, *First-Price Path Auctions*, EC 2005, 203-212.
- [20] A. Iwasaki, D. Kempe, Y. Saito, M. Salek, M. Yokoo, *False-Name-Proof Mechanisms for Hiring a Team*, WINE 2007, 245-256.
- [21] A. R. Karlin, D. Kempe, T. Tamir, *Beyond VCG: Frugality of Truthful Mechanisms*, FOCS 2005, 615-626.
- [22] D. Kempe, personal communication.
- [23] D. Kempe, M. Salek, C. Moore, *Frugal and Truthful Auctions for Vertex Covers, Flows, and Cuts*, arXiv:0912.3310.

- [24] V. Krishna, *Auction Theory*, Academic Press, 2002.
- [25] C. L. Li, S. T. McCormick, D. Simchi-Levi, *The Complexity of Finding Two Disjoint Paths with min-max Objective Function*, Discrete Applied Mathematics, V.26(1), 105-115, 1990.
- [26] N. Nisan, A. Ronen, *Algorithmic Mechanism Design*, STOC 1999, 129-140.
- [27] K. Talwar, *The Price of Truth: Frugality in Truthful Mechanisms*, STACS 2003, 608-619.
- [28] W. Vickrey, *Counterspeculation, Auctions and Competitive Sealed Tenders*, Journal of Finance, V.16, 8-37, 1961.
- [29] Q. Yan, *On the Price of Truthfulness in Path Auctions*, WINE 2007, 584-589.

A Proof of Theorem 3.1

Proof. For any agent $e \in \mathcal{E}$ and given bids of other agents, we will analyze the utility of e in terms of his bid. There are the following two cases.

Case 1. Agent e is not dropped out in the pruning process when bidding $b(e) = c(e)$, i.e. $e \in \mathcal{E}^*$. By the definition of $t_1(e)$, we know that $t_1(e) \geq c(e)$. Consider the situation where e bids another value $b'(e) \neq b(e)$. If $b'(e) > t_1(e)$, then $e \notin \mathcal{E}^*$ and his utility is 0. If $b'(e) \leq t_1(e)$, by the bid-independence property, we know that the subset \mathcal{E}^* remains the same. Given this fact, the structure of the graph \mathcal{H} does not change, which implies that the eigenvectors and eigenvalues of its adjacency matrix do not change either. Hence, the threshold value $t_2(e)$ will not change, which implies that the payment to agent e , $p(e) = \min\{t_1(e), t_2(e)\}$, will not change.

Case 2. Agent e is dropped out in the pruning process when bidding $b(e) = c(e)$, i.e. $e \notin \mathcal{E}^*$. Consider the situation where e bids another value $b'(e) \neq b(e)$ and is not dropped out. By monotonicity and bid-independence, we know that $b'(e) \leq t_1(e) \leq b(e) = c(e)$. Hence, even though e could be a winner by bidding $b'(e)$, his payment is at most $t_1(e) \leq c(e)$, which implies that he cannot obtain a positive utility.

From the above two cases, we know that the utility of each agent is maximized by bidding his true cost, and hence the mechanism is truthful. \square

B Analysis of PRUNING-LIFTING k -PATHS MECHANISM

B.1 Lower Bound on ν

Our analysis of PRUNING-LIFTING k -PATHS MECHANISM relies on the characterization of Nash flows presented in [17].

Theorem B.1 ([17]). *Let $G = (V, E)$ be a directed graph with weight $w(e)$ on each edge $e \in E$. Given two specific vertices $s, t \in V$, assume that there are k edge-disjoint paths from s to t . Let P_1, P_2, \dots, P_k be such k edge-disjoint s - t paths so that its total weight $L \triangleq \sum_{i=1}^k w(P_i)$ is minimized, where $w(P_i) = \sum_{e \in P_i} w(e)$. Further, it is known that for every edge $e \in E$, the graph $G - \{e\}$ has k edge-disjoint s - t paths with the same total weight L . Then essentially there exist $k + 1$ edge-disjoint s - t paths in G such that each of them is a shortest path from s to t .*

We will first show how Theorem B.1 implies Lemma 5.3. For completeness, we will then present a proof of Theorem B.1; this proof also appears in [17].

Proof of Lemma 5.3. Fix a cost vector \mathbf{c} . Let E' be the winning set with respect to \mathbf{c} , and consider a bid vector \mathbf{b} that satisfies conditions (1)–(3) in the definition of $\nu(\mathbf{c})$. Let $p(\mathbf{b})$ denote the total payment under \mathbf{b} . The set E' contains k edge-disjoint s - t paths. By condition (2), no agent in E' can obtain more revenue by increasing his bid. That is, for any $e \in E'$, there are k edge-disjoint s - t paths in $G \setminus \{e\}$ with the same total bid as E' . Applying Theorem B.1 with $w(e) = b(e)$, we obtain that there are $k + 1$ edge-disjoint shortest s - t paths with length $\frac{p(\mathbf{b})}{k}$ each w.r.t \mathbf{b} . Consider the subgraph G' composed by these $k + 1$ edge-disjoint paths. We know that $\delta_{k+1}(G', \mathbf{c}) \leq \frac{p(\mathbf{b})}{k}$ as $b(e) \geq c(e)$ for any edge e , i.e. the longest s - t path in G' w.r.t to \mathbf{c} is at most $\frac{p(\mathbf{b})}{k}$. Hence,

$$p(\mathbf{b}) \geq k \cdot \delta_{k+1}(G', \mathbf{c}) \geq k \cdot \delta_{k+1}(G, \mathbf{c}).$$

As this holds for any vector \mathbf{b} that satisfies conditions (1)–(3), it follows that $\nu(\mathbf{c}) \geq \delta_{k+1}(G, \mathbf{c})$. \square

In the rest of this section, we prove Theorem B.1.

Proof of Theorem B.1. Given the graph G and integer k , we construct a flow network $\mathcal{N}_k(G)$ as follows: We introduce two extra nodes s_0 and t_0 and two edges s_0s and tt_0 . The set of vertices of $\mathcal{N}_k(G)$ is $V \cup \{s_0, t_0\}$ and the set of edges is $E \cup \{s_0s, tt_0\}$. The capacity $\text{cap}(\cdot)$ and cost per bulk capacity $\text{cost}(\cdot)$ for each edge in $\mathcal{N}_k(G)$ is defined as follows:

- $\text{cap}(s_0s) = \text{cap}(tt_0) = k$ and $\text{cost}(s_0s) = \text{cost}(tt_0) = 0$.
- $\text{cap}(e) = 1$ and $\text{cost}(e) = w(e)$, for $e \in E$.

Given the above construction, every path from s to t in G naturally corresponds to a bulk flow from s_0 to t_0 in $\mathcal{N}_k(G)$. Hence, the set of k edge-disjoint paths P_1, P_2, \dots, P_k in G corresponds to a flow \mathcal{F}_G of size k in $\mathcal{N}_k(G)$. In addition, the minimality of $L = \sum_{i=1}^k w(P_i)$ implies that \mathcal{F}_G achieves the minimum cost (which is L) for all integer-valued flows of size k , i.e. maximum flow in $\mathcal{N}_k(G)$. Since all capacities of $\mathcal{N}_k(G)$ are integers, we can conclude that \mathcal{F}_G has the minimum cost among all real maximum flows in $\mathcal{N}_k(G)$. (The proof of this fact can be found in, e.g., [16].)

For simplicity, we denote the subgraph $G - \{e\}$ by $G - e$. By the fact that for any $e \in E$, the subgraph $G - e$ has k edge-disjoint s - t paths with the same total weight L , we know that in the network $\mathcal{N}_k(G - e)$, there still is an integer-valued flow \mathcal{F}_{G-e} of size k and cost L . So \mathcal{F}_{G-e} is also an integer-valued flow of size k and cost L in $\mathcal{N}_k(G)$. Define a real-valued flow in $\mathcal{N}_k(G)$ by $\mathcal{F} = \frac{1}{|E|} \sum_{e \in E} \mathcal{F}_{G-e}$. We have the following observations:

1. It is clear that $\mathcal{F}(e) \leq \text{cap}(e)$ for every arc $e \in \mathcal{N}_k(G)$, where $\mathcal{F}(e)$ is the amount of flow on edge e in \mathcal{F} , as we have taken the arithmetic average of the flows in the network.
2. \mathcal{F} has cost $\frac{1}{|E|} \sum_{e \in E} \text{cost}(\mathcal{F}_{G-e}) = \frac{1}{|E|} \cdot |E| \cdot L = L$.
3. Since $\mathcal{F}_{G-e}(s_0s) = k$ for any $e \in E$, we have $\mathcal{F}(s_0s) = k$. In addition, as each \mathcal{F}_{G-e} is a feasible flow that satisfies all conservation conditions and \mathcal{F} is defined by the arithmetic average of all \mathcal{F}_{G-e} 's, we know that \mathcal{F} also satisfies all conservation conditions.

Therefore, \mathcal{F} is a minimum cost maximum flow in $\mathcal{N}_k(G)$. In addition, \mathcal{F} has the following nice property, which plays a fundamental role for the proof:

- For every edge $e \in \mathcal{N}_k(G)$ except s_0s and tt_0 , we have $\mathcal{F}(e) \leq \text{cap}(e) - \frac{1}{|E|}$, as \mathcal{F}_{G-e} does not flow through e , i.e. $\mathcal{F}_{G-e}(e) = 0$, and $\mathcal{F}_{G-e'}(e)$ is either 0 or 1 for any $e' \in E$.

Let $E_+ = \{e \in \mathcal{N}_k(G) \mid \mathcal{F}(e) > 0\}$. Suppose that there is a path $P' = (e_1, e_2, \dots, e_r)$ from s_0 to t_0 which goes only along arcs in E_+ and is not a shortest path w.r.t $\text{cost}(\cdot)$ from s_0 to t_0 in $\mathcal{N}_k(G)$. Let $\epsilon = \min \left\{ \mathcal{F}(e_1), \mathcal{F}(e_2), \dots, \mathcal{F}(e_r), \frac{1}{|E|} \right\}$. Since $P' \subseteq E_+$, we have $\epsilon > 0$. Let P be a shortest path w.r.t $\text{cost}(\cdot)$ from s_0 to t_0 in $\mathcal{N}_k(G)$. Define a new flow \mathcal{F}' from \mathcal{F} by adding ϵ amount of flow on path P and removing ϵ amount of flow from path P' . We have the following observations about \mathcal{F}' :

1. The size of flow \mathcal{F}' is k .
2. \mathcal{F}' satisfies all conservation conditions as it is a linear combination of three flows from s_0 to t_0 .
3. By the definition of ϵ , the amount of flow of each edge is non-negative in \mathcal{F}' . Further, \mathcal{F}' satisfies the capacity constraints. This follows from the facts that $\epsilon \leq \frac{1}{|E|}$ and the above property established for \mathcal{F} .

4. The cost of \mathcal{F}' is smaller than L because $\text{cost}(\mathcal{F}') = \text{cost}(\mathcal{F}) - \epsilon(\text{cost}(P') - \text{cost}(P))$, which is smaller than $L = \text{cost}(\mathcal{F})$ as $\text{cost}(P) < \text{cost}(P')$ by the assumption.

Hence, \mathcal{F}' is a flow of size k in $\mathcal{N}_k(G)$ with cost smaller than \mathcal{F} , a contradiction. Thus, every path from s_0 to t_0 in $\mathcal{N}_k(G)$ along the edges in E_+ is a shortest path w.r.t $\text{cost}(\cdot)$.

Consider a new network $\mathcal{N}'_{k+1}(G)$ obtained from $\mathcal{N}_{k+1}(G)$ by restricting edges in E_+ . (Note that the only difference between $\mathcal{N}_{k+1}(G)$ and $\mathcal{N}_k(G)$ is the capacity on edges s_0s and tt_0 is $k+1$ rather than k .) We claim that in this network there is an integer-valued flow of size $k+1$. Suppose otherwise, by Ford-Fulkerson Theorem [13], there is a cut (S_{s_0}, T_{t_0}) in $\mathcal{N}'_{k+1}(G)$ with size less than or equal to k . By definition, in $\mathcal{N}'_{k+1}(G)$ we have $\text{cap}(s_0s) = k+1$ and $\text{cap}(tt_0) = k+1$, which implies that $s_0, s \in S_{s_0}$ and $t_0, t \in T_{t_0}$. By the definition of E_+ , we know that the total amount of flows of \mathcal{F} on the cut (S_{s_0}, T_{t_0}) is k . Since $\mathcal{F}(e) < 1$ for any edge e , we can conclude that there are at least $k+1$ edges from S_{s_0} to T_{t_0} in E_+ . This leads to a contradiction, because we have showed that the size of the cut (S_{s_0}, T_{t_0}) is less than or equal to k .

Therefore, we can find an integer-valued flow of size $k+1$ on edges in E_+ in the network $\mathcal{N}_{k+1}(G)$. Such a flow can be thought as the union of $k+1$ edge-disjoint paths from s_0 to t_0 . We know that every such path going along edges in E_+ is a shortest path from s_0 to t_0 . This in turn concludes the proof, since we have found $k+1$ edge-disjoint shortest paths from s to t in G . \square




Figure 1: An instance of 1-path problem on which PRUNING-LIFTING k -PATHS MECHANISM does not choose the path that minimizes $\delta_2(G, \mathbf{c})$

B.2 Upper Bound on the Payment of PRUNING-LIFTING k -PATHS MECHANISM

In this subsection, we will prove Lemma 5.4. Fix a cost vector \mathbf{c} and set $G^* = G^*(\mathbf{c})$, $\mathcal{H} = \mathcal{H}(G^*(\mathbf{c}))$, $\alpha = \alpha(G^*(\mathbf{c}))$. Observe that since G^* is the cheapest collection of k edge-disjoint paths in G^* , it is necessarily cycle-free. We say that a vertex v is an *articulation point* for G^* if it lies on any s - t path in G . Suppose that G^* has $a+1$ *articulation points* $s = v_1, v_2, \dots, v_{a+1} = t$ that subdivide G^* into a parts G_1^*, \dots, G_a^* . Denote by \mathcal{H}_j the induced subgraph of \mathcal{H} with vertex set $\{v_e \mid e \in G_j^*\}$, $j = 1, \dots, a$. For any graph Γ , let $V(\Gamma)$ and $E(\Gamma)$ denote the sets of vertices and edges of Γ , respectively. The following intermediate lemma describes the structure of the graph \mathcal{H} .

Lemma B.2. *Graph \mathcal{H} is connected if and only if G^* has no articulation points. Hence, for each j , \mathcal{H}_j is connected.*

For each vertex $v \in V(\mathcal{H})$, let e_v be the corresponding edge in G^* . Before we prove Lemma B.2, we need the following two observations.




Figure 2: An example of the construction of \mathcal{H} from G^* for $k = 2$

Claim B.3. *Let u and v be two vertices of \mathcal{H} , then $uv \notin E(\mathcal{H})$ if and only if there is an s - t path in G^* going through both e_u and e_v .*

Proof. If there is a path $P \subseteq G^*$ such that $e_u, e_v \in P$, then in $G^* \setminus \{e_u, e_v\}$ there are k edge-disjoint s - t paths. Hence there is no edge between u and v . Conversely, if $uv \notin E(\mathcal{H})$, then $G^* \setminus \{e_u, e_v\}$ has k edge-disjoint s - t paths. Removing these k paths from G^* leads to an s - t path going through e_u and e_v . \square

Claim B.4. *Let P be an s - t path in G^* , corresponding to the vertex sequence v_1, \dots, v_ℓ of \mathcal{H} , and let v be another vertex of \mathcal{H} . Then the following holds:*

1. *There are integral i and j , $1 \leq i \leq j \leq \ell$, such that $vv_r \in E(\mathcal{H})$ for $i \leq r \leq j$, and $vv_r \notin E(\mathcal{H})$ for $1 \leq r < i$ and $j < r \leq \ell$.*
2. *Let $\mathcal{L}(v) = \{v_r \mid 1 \leq r \leq \ell, vv_r \notin E(\mathcal{H})\}$, then in G^* there is a s - t path containing all vertices of $\mathcal{L}(v)$ and v .*

Proof.

1. Since $G^* \setminus \{e_v\}$ contains only k edge-disjoint s - t paths, there is a cut of size $k + 1$ in G^* containing e_v and this cut must contain an edge of P . Thus, there exists a vertex v_i with $1 \leq i \leq \ell$ adjacent to v in \mathcal{H} . For the first part of the claim, it remains to show that if v is adjacent to v_p and v_q with $p < q$, then so is v_r for any r with $p \leq r \leq q$. Suppose otherwise, then by Claim B.3 there should be a path P' going through e_v and e_{v_r} for some $p < r < q$. If P' goes through e_v earlier than e_{v_r} , we can construct a path in G^* going through both e_v and e_{v_q} . Indeed, we can concatenate the prefix of P' that ends with e_{v_p} and the suffix of P that starts with e_{v_p} ; this path is simple since G^* is cycle-free. By a similar argument, if P' goes through e_{v_r} earlier than e_v , then we can construct a path in G^* going through e_{v_p} and e_v . For both cases we arrive at a contradiction with Claim B.3, since v_p and v_q are adjacent to v in \mathcal{H} .
2. This part can be derived from the first part and Claim B.3 applied to two pairs of vertices v_{i-1}, v and v_{j+1}, v . \square

We are now ready to prove Lemma B.2.

Proof of Lemma B.2. Suppose that G^* contains an articulation point $v \in V(G^*)$. Then we can split G^* into the union of two subgraphs G_1 (from s to v) and G_2 (from v to t). By the definition of \mathcal{H} , it can be seen that for any $e_1 \in E(G_1)$ and $e_2 \in E(G_2)$, $v_{e_1}v_{e_2} \notin E(\mathcal{H})$. Hence, if \mathcal{H} is connected, then G^* has no articulation points.

On the other hand, suppose that G^* has no articulation points. We can endow the set of nodes of G^* with a partial order by putting $u > v$ if there is a path in G^* from u to v . Suppose that \mathcal{H} is disconnected. Let C be a connected component of \mathcal{H} . Let $S(C)$ be the set of endpoints of the edges in G^* corresponding to those vertices in C . We claim that $S(C)$ has a unique smallest element, denoted by $\rho(C)$. Suppose otherwise that $S(C)$ has at least two different smallest elements v_1 and v_2 with $e_1 = u_1v_1 \in E(G^*)$ and $e_2 = u_2v_2 \in E(G^*)$. Assume without loss of generality that $v_1 \neq t$. Then there is $v_3 \in V(G^*)$ such that $e_3 = v_1v_3 \in E(G^*)$. There could not be a path going through both e_2 and e_3 , since in this case either $v_2 > v_1$ or $v_1 > v_2$. Hence e_2 and e_3 are adjacent in \mathcal{H} , which implies that $e_3 \in C$. But then we have $v_3 \in S(C)$ and $v_3 < v_1$, a contradiction to the assumption that v_1 is a smallest element.

Since $e_1 = uv \in E(G^*)$ and $e_2 = wv \in E(G^*)$ are necessarily adjacent in \mathcal{H} , two different connected components C_1 and C_2 of \mathcal{H} cannot have the same smallest element, i.e. $\rho(C_1) \neq \rho(C_2)$. Let us take the component C with the smallest element $\rho(C) \neq t$. It is also clear by the definition of $\rho(C)$ that $\rho(C) \neq s$. Let e be a vertex of C with the ending point at $\rho(C)$ in G^* . We want to show that $\rho(C)$ is an articulation point for G^* . Suppose otherwise that there is a path P in G^* which does not go through $\rho(C)$. Let $v_e \in \mathcal{H}$ correspond to e and $v_1, \dots, v_\ell \in \mathcal{H}$ correspond to P . Applying Claim B.4 to P and v_e , we get that v_j is adjacent to v_e in \mathcal{H} and there is a path from endpoint of v_e in G^* , i.e. $\rho(C)$, to the endpoint of v_j . Thus we arrive at a contradiction, since $v_j \in C$ and endpoint of v_j is smaller than $\rho(C)$. \square

We will now prove Lemma 5.4.

Proof of Lemma 5.4. Suppose first that G^* has no articulation points. Let P^* be a losing path in G^* according to the mechanism. Then P^* is the most expensive path in G^* w.r.t. scaled cost vector. Assume that P^* is represented as a sequence of vertices v_1, v_2, \dots, v_ℓ in \mathcal{H} and let v be another vertex of \mathcal{H} . Let $i(v)$ and $j(v)$ be the two integers, $1 \leq i(v) \leq j(v) \leq \ell$, as defined in Claim B.4. Note that the threshold bid of v is at most $\sum_{r=i(v)}^{j(v)} \frac{c(v_r)}{w(v_r)} w(v)$, since otherwise by Claim B.4, one can find a more expensive path (w.r.t. scaled costs) going through $\mathcal{L}(v)$ and v than P^* . Thus we have the following upper bound on the total payment of the mechanism:

$$\begin{aligned} \sum_{v \notin P^*} \sum_{r=i(v)}^{j(v)} \frac{c(v_r)}{w(v_r)} w(v) &= \sum_{r=1}^{\ell} \sum_{v \in N_{\mathcal{H}}(v_r)} \frac{c(v_r)}{w(v_r)} w(v) \\ &= \sum_{r=1}^{\ell} \frac{c(v_r)}{w(v_r)} \sum_{v \in N_{\mathcal{H}}(v_r)} w(v) \\ &= \sum_{r=1}^{\ell} \frac{c(v_r)}{w(v_r)} \alpha(G^*) w(v_r) \\ &= \alpha(G^*) \sum_{t=1}^{\ell} c(v_t). \end{aligned}$$

As $L(G, \mathbf{c}) \geq \sum_{t=1}^{\ell} c(v_t)$, this proves the lemma for this case.

For the general case where there are articulation points in G^* , by using the same analysis on each connected component \mathcal{H}_i of \mathcal{H} , we can upper-bound the total payment of our mechanism by $\sum_{\mathcal{H}_i} \alpha_i \sum_{r=1}^{\ell_i} c(v_{i,r})$, where $v_{i,1}, v_{i,2}, \dots, v_{i,\ell_i}$ are a sequence of vertices in each \mathcal{H}_i corresponding to the most expensive path P^* , and α_i is the maximum eigenvalue of the adjacency matrix of \mathcal{H}_i . As $\alpha(G^*) = \max_i \alpha_i$, and the most expensive path in G^* is a concatenation of the most expensive paths in G_1^*, \dots, G_a^* , the lemma is proven. \square

B.3 Proof of Theorem 5.1

We will now show how Lemmas 5.3 and 5.4 imply Theorem 5.1.

Proof of Theorem 5.1. Fix an arbitrary cost vector \mathbf{c} . Suppose that in the pruning stage we pick a graph G^* . By Lemma 5.4, the total payment of our mechanism is at most $\alpha(G^*)L(G, \mathbf{c})$. Since $G^* \in \mathcal{G}^{k+1}$, we have $\alpha(G^*) \leq \alpha_{k+1}$. Consider a collection P_1, \dots, P_{k+1} of $k+1$ edge-disjoint paths in G such that $\delta_{k+1}(P_1, \dots, P_{k+1}, \mathbf{c}) = \delta_{k+1}(G, \mathbf{c})$. Since G^* is the cheapest collection of $k+1$ edge-disjoint paths in G , we have $\sum_{e \in G^*} c(e) \leq \sum_{i=1}^{k+1} \sum_{e \in P_i} c(e)$. We obtain

$$L(G, \mathbf{c}) \leq \sum_{e \in G^*} c(e) \leq \sum_{i=1}^{k+1} \sum_{e \in P_i} c(e) \leq (k+1)\delta_{k+1}(P_1, \dots, P_{k+1}, \mathbf{c}) = (k+1)\delta_{k+1}(G, \mathbf{c}),$$

where the last inequality follows from the definition of $\delta_{k+1}(P_1, \dots, P_{k+1}, \mathbf{c})$. Thus, the frugality ratio of PRUNING-LIFTING k -PATHS MECHANISM on \mathbf{c} is at most

$$\frac{\alpha(G^*)L(G, \mathbf{c})}{k\delta_{k+1}(G, \mathbf{c})} \leq \frac{\alpha_{k+1}(k+1)\delta_{k+1}(G, \mathbf{c})}{k\delta_{k+1}(G, \mathbf{c})} = \frac{\alpha_{k+1}(k+1)}{k}.$$

\square

C μ -Frugality Analysis: Proof of Theorem 5.6

To prove Theorem 5.6, we need to show that we can lower-bound $\mu(\mathbf{c})$ in terms of $L(G, \mathbf{c})$.

Consider an arbitrary network \mathcal{F} with source s , sink t , integer edge capacities and costs per unit flow that are given by a vector \mathbf{c} . Let M be the size of the maximum flow in \mathcal{F} . For any (real) $x \in [0, M]$, let $C(x)$ be the cost of a cheapest flow of size x in \mathcal{F} (i.e. the sum of costs on all edges, where the cost on an edge e is the amount of flow times $\mathbf{c}(e)$). The following lemma establishes several properties of the function $C(x)$ that will be used in our proof.

Lemma C.1.

1. $C(x)$ is a convex function on $[0, M]$.
2. For any integer $i \leq M-1$, $C(x)$ is a linear function on the interval $[i, i+1]$.

Proof.

- **Convexity.** It suffices to show that for any $0 \leq \alpha \leq 1$ we have

$$\alpha C(x) + (1-\alpha)C(y) \geq C(\alpha x + (1-\alpha)y).$$

Let f_x and f_y be cheapest flows of size x and y , respectively. Clearly, their respective costs are $C(x)$ and $C(y)$. Then the flow $f = \alpha f_x + (1 - \alpha)f_y$ is a flow of size $(\alpha x + (1 - \alpha)y)$ and cost $\alpha C(x) + (1 - \alpha)C(y)$ that satisfies all capacity constraints. Thus,

$$\alpha C(x) + (1 - \alpha)C(y) = \mathbf{c}(f) \geq C(\alpha x + (1 - \alpha)y).$$

- **Linearity on intervals.** We first show that $C(x)$ is linear on the interval $[0, 1]$. Let us fix $x_0 \in [0, 1]$ and let f be a cheapest flow of size x_0 . We can represent f as a finite sum of positive flows along s - t paths p_1, \dots, p_l , i.e.,

$$f = \sum_{i=1}^l \varepsilon_i p_i.$$

We know that $C(1)$ is the cost of cheapest path p . Thus we have

$$C(x_0) = \mathbf{c}(f) \geq \mathbf{c}(p) \sum_{i=1}^l \varepsilon_i = C(1)x_0.$$

On the other hand, since $C(x)$ is convex, we have $x_0 C(1) + (1 - x_0)C(0) \geq C(x_0)$. Hence $C(x_0) = x_0 C(1)$.

In general, for the interval $[i, i + 1]$ we first take a cheapest i -flow f_i (which we can choose to be integer) and consider the residual network $\mathcal{F}_i = \mathcal{F} - f_i$. We can then apply the argument for the $[0, 1]$ case to \mathcal{F}_i .

□

Our next lemma bounds $\mu(\mathbf{c})$ in terms of the difference between the cost of the cheapest flow of size k and that of the cheapest flow of size $k + 1$.

Lemma C.2. *Let $(\mathcal{E}, \mathcal{F})$ be a k -path system given by a directed graph $G = (V, E)$, source s and sink t , and let \mathbf{c} be its cost vector. Then for the function $C(x)$ defined above we have $k \cdot (C(k + 1) - C(k)) \leq \mu(\mathbf{c})$.*

Proof. For any cost vector $\mathbf{y} \in \mathbb{R}^{|E|}$, let $C_{\mathbf{y}}(x)$ denote the cost of the cheapest flow of size x in G with respect to the cost vector \mathbf{y} ; we have $C_{\mathbf{c}}(x) = C(x)$.

Let f_k be a cheapest flow of size k , and let f_{k+1} be a cheapest flow of size $k + 1$, both with respect to cost vector \mathbf{c} . Let n_k denote the number of edges in f_k . Assume without loss of generality that the edges in f_k are labeled as e_1, \dots, e_{n_k} .

We will now gradually increase the costs of edges in f_k so that the resulting cost vector \mathbf{y} satisfies certain conditions. Specifically, we start with $\mathbf{y} = \mathbf{c}$. Then, at i -th step, $i = 1, \dots, n_k$, we increase $\mathbf{y}(e_i)$ as much as possible subject to the following constraints:

- (a) $\mathbf{y}(f_k) = \sum_{e \in f_k} \mathbf{y}(e) = C_{\mathbf{y}}(k)$, i.e., f_k must remain the cheapest k -flow w.r.t. cost vector \mathbf{y} .
- (b) $C_{\mathbf{y}}(k + 1) - C_{\mathbf{y}}(k) = C(k + 1) - C(k)$, i.e. $C_{\mathbf{y}}(k + 1) - C_{\mathbf{y}}(k)$ does not change.

Since our k -path system is monopoly-free, in the end, all entries of \mathbf{y} are finite. Further, it is not hard to see that when the process is over, we cannot increase the cost of any edge in f_k without violating (a) or (b).

Now, for each edge $e \in f_k$, we will define the *tight flow* $f(e)$ as below to be a flow that prevented us from raising $\mathbf{y}(e)$ beyond its current value. Specifically, consider each edge $e_i \in f_k$. Suppose first

that when we were raising $\mathbf{y}(e_i)$, we had to stop because constraint (a) became tight. In this case, let $f(e_i)$ be some cheapest flow of size k in $G \setminus \{e_i\}$ with respect to the costs \mathbf{y} at the end of stage i . Now, suppose that when we were raising $\mathbf{y}(e_i)$, constraint (b) became tight first. In this case, let $f(e_i)$ be some cheapest flow of size $k + 1$ in $G \setminus \{e_i\}$ with respect to the costs \mathbf{y} at the end of stage i . Observe that f_k remains a cheapest k -flow throughout the process; further, for all $e \in f_k$, the flow $f(e)$ is a cheapest flow of its size in G with respect to the final cost vector as well. In the following we consider the cost vector \mathbf{y} at the end of the process.




Figure 3: The graph of $C_{\mathbf{y}}(x)$

Let f^* be the average of all tight flows, i.e., set

$$f^* = \frac{1}{n_k} \sum_{e \in f_k} f(e).$$

Let q be the size of f^* ; we have $k \leq q \leq k + 1$. Note that f^* is a cheapest flow of size q by the second statement of lemma C.1, as it is a convex combination of cheapest flows of size k and cheapest flows of size $k + 1$. Further, since $e \notin f(e)$ for any $e \in f_k$, the amount of flow that passes through each edge e in f^* is strictly less than 1. Thus, for a sufficiently small $\epsilon > 0$, flow $f^* + \epsilon f_k$ is a valid flow of size $q + \epsilon k$ in G . Moreover, we have $C_{\mathbf{y}}(q + \epsilon k) \leq \mathbf{y}(f^* + \epsilon f_k) = C_{\mathbf{y}}(q) + \epsilon C_{\mathbf{y}}(k)$.

This observation, together with the convexity of $C_{\mathbf{y}}(x)$, allows us to derive that $C_{\mathbf{y}}(x)$ is a linear function on the interval $[0, k + 1]$. Indeed, by convexity of $C_{\mathbf{y}}(x)$ we have

$$\begin{aligned} C_{\mathbf{y}}(k) &\leq \frac{q + (\epsilon - 1)k}{q + \epsilon k} C_{\mathbf{y}}(0) + \frac{k}{q + \epsilon k} C_{\mathbf{y}}(q + \epsilon k) = \frac{k}{q + \epsilon k} C_{\mathbf{y}}(q + \epsilon k) \\ C_{\mathbf{y}}(q) &\leq \frac{\epsilon k}{q + \epsilon k} C_{\mathbf{y}}(0) + \frac{q}{q + \epsilon k} C_{\mathbf{y}}(q + \epsilon k) = \frac{q}{q + \epsilon k} C_{\mathbf{y}}(q + \epsilon k) \end{aligned}$$

If any of the two inequalities above is an equality, then $C_{\mathbf{y}}(x)$ is linear on $[0, k + 1]$ and we are done. Otherwise, both of these inequalities are strict, and we can write

$$C_{\mathbf{y}}(q + \epsilon k) \leq C_{\mathbf{y}}(q) + \epsilon C_{\mathbf{y}}(k) < \frac{q}{q + \epsilon k} C_{\mathbf{y}}(q + \epsilon k) + \epsilon \frac{k}{q + \epsilon k} C_{\mathbf{y}}(q + \epsilon k) = C_{\mathbf{y}}(q + \epsilon k).$$

The contradiction shows that $C_{\mathbf{y}}(x)$ is a linear function on $[0, k + 1]$.

Since \mathbf{y} satisfies conditions (1) and (2) in the definition of $\mu(\mathbf{c})$, we obtain

$$\mu(\mathbf{c}) \geq C_{\mathbf{y}}(k) = k(C_{\mathbf{y}}(k + 1) - C_{\mathbf{y}}(k)) = k(C_{\mathbf{c}}(k + 1) - C_{\mathbf{c}}(k)),$$

where the first equality follows from linearity of $C_{\mathbf{y}}(x)$ and last equality holds by construction of \mathbf{y} . Thus, the lemma is proven. \square

Recall now that $L(G, c)$ by definition is the longest path in $G^*(\mathbf{c})$, where $G^* = G^*(\mathbf{c})$ is the cheapest $k + 1$ -flow w.r.t. \mathbf{c} . Hence, $C(k + 1) = \mathbf{c}(G^*)$. Further we can decompose G^* into the sum of a path with the cost $L(G, c)$ and some k -flow, which implies that $C(k + 1) \geq L(G, c) + C(k)$. Rewriting the last inequality we get $C(k + 1) - C(k) \geq L(G, c)$. Combining this observation with Lemma 5.4 we easily derive Theorem 5.6.

D Proof of Lemma 6.3

Let C_v be some clique of size ρ_v in \mathcal{CL}_v , and consider the bid vector \mathbf{b} given by $b(u) = x$ if $u \in C_v$ and $b(u) = 0$ if $u \in V \setminus C_v$. Since C_v is a clique, any vertex cover for G must contain at least $\rho_v - 1$ vertices of C_v . Thus, any cheapest feasible set with respect to the true costs contains all vertices in $C_v \setminus \{v\}$; let S denote some such set. Moreover, for any $u \in C_v \setminus \{v\}$, any vertex cover that does not contain u must contain v , so \mathbf{b} satisfies condition (2) with respect to the set S in the definition of the benchmark ν . To see that \mathbf{b} also satisfies condition (3), note that if any vertex in $C_v \setminus \{v\}$ decides to raise its bid, it can be replaced by its neighbors at cost x . Now, consider any $w \in (V \setminus C_v) \cap S$. The vertex w cannot be adjacent to all vertices in C_v , since otherwise C_v would not be a maximal clique. Thus, if $w \in S$, we can obtain a vertex cover of cost $x(\rho_v - 1)$ that does not include w by taking all vertices of cost 0 as well as all vertices in C_v that are adjacent to w .